

Machine Learning

Volume 22, Nos. 1/2/3, January/February/March 1996

Special Issue on Reinforcement Learning

Guest Editor: Leslie Pack Kaelbling

Editorial	Thomas G. Dietterich	5
Introduction	Leslie Pack Kaelbling	7
Efficient Reinforcement Learning through Symbiotic Evolution	David E. Moriarty and Risto Miikkulainen	11
Linear Least-Squares Algorithms for Temporal Difference Learning	Steven J. Bradtke and Andrew G. Barto	33
Feature-Based Methods for Large Scale Dynamic Programming	John N. Tsitsiklis and Benjamin Van Roy	59
On the Worst-Case Analysis of Temporal-Difference Learning Algorithms	Robert E. Schapire and Manfred K. Warmuth	95
Reinforcement Learning with Replacing Eligibility Traces	Satinder P. Singh and Richard S. Sutton	123
Average Reward Reinforcement Learning: Foundations, Algorithms, and Empirical Results	Sridhar Mahadevan	159
The Loss from Imperfect Value Functions in Expectation-Based and Minimax-Based Tasks	Matthias Heger	197
The Effect of Representation and Knowledge on Goal-Directed Exploration with Reinforcement-Learning Algorithms	Sven Koenig and Reid G. Simmons	227
Creating Advice-Taking Reinforcement Learners	Richard Maclin and Jude W. Shavlik	251
Technical Note: Incremental Multi-Step Q-Learning	Jing Peng and Ronald J. Williams	283

Volume 23, No. 1, April 1996

Representation of Finite State Automata in Recurrent Radial Basis Function Networks	Paolo Frasconi, Marco Gori, Marco Maggini and Giovanni Soda	5
Scaling Up Inductive Learning with Massive Parallelism	Foster J. Provost and John M. Aronis	33
Classification by Feature Partitioning	H. Altay Güvenir and Izzet Sirin	47
Learning in the Presence of Concept Drift and Hidden Contexts	Gerhard Widmer and Miroslav Kubat	69

Book Review: <i>Inductive Logic Programming: Techniques and Applications</i>	Michael Pazzani	103
A Reply to Pazzani's Book Review of <i>Inductive Logic Programming: Techniques and Applications</i>	Nada Lavrač and Sašo and Džeroski	109

Volume 23, Nos. 2/3, May/June 1996

Special Issue on Robot Learning

Guest Editors: Judy A. Franklin, Tom M. Mitchell and Sebastian Thrun

Introduction	Judy A. Franklin, Tom M. Mitchell and Sebastian Thrun	117
Real-World Robotics: Learning to Plan for Robust Execution	Scott W. Bennett and Gerald F. DeJong	121
Robot Programming by Demonstration (RPD): Supporting the Induction by Human Interaction	H. Friedrich, S. Münch, R. Dillmann, S. Bocionek and M. Sassin	163
Performance Improvement of Robot Continuous-Path Operation through Iterative Learning Using Neural Networks	Peter C.Y. Chen, James K. Mills and Kenneth C. Smith	191
Learning Controllers for Industrial Robots	C. Baroglio, A. Giordana, M. Kaiser, M. Nuttin and R. Piola	221
Active Learning for Vision-Based Robot Grasping	Marcos Salganicoff, Lyle H. Ungar and Ruzena Bajcsy	251
Purposive Behavior Acquisition for a Real Robot by Vision-Based Reinforcement Learning	Minoru Asada, Shoichi Noda, Sukoya Tawaratsumida and Koh Hosoda	279
Learning Concepts from Sensor Data of a Mobile Robot	Volker Klingspor, Katharina J. Morik and Anke D. Rieger	305

Volume 24, No. 1, July 1996

BEXA: A Covering Algorithm for Learning Propositional Concept Descriptions	Hendrik Theron and Ian Cloete	5
Technical Note: Some Properties of Splitting Criteria	Leo Breiman	41
Stacked Regressions	Leo Breiman	49
On Learning Visual Concepts and DNF Formulae	Eyal Kushilevitz and Dan Roth	65

Volume 24, No. 2, August 1996

Editorial	Thomas G. Dietterich	91
A Lattice Conceptual Clustering System and Its Application to Browsing Retrieval	Claudio Carpineto and Giovanni Romano	95

Bagging Predictors	<i>Leo Breiman</i>	123
Unifying Instance-Based and Rule-Based Induction	<i>Pedro Domingos</i>	141

Volume 24, No. 3, September 1996

Error Reduction through Learning Multiple Descriptions	<i>Kamal M. Ali and Michael J. Pazzani</i>	173
A Decision-Tree Model of Balance Scale Development	<i>William C. Schmidt and Charles X. Ling</i>	000
Efficient Incremental Induction of Decision Trees	<i>Dimitrios Kalles and Tim Morris</i>	000

Volume 25, No. 1, October 1996

Exploration Bonuses and Dual Control.... <i>Peter Dayan and Terrence J. Sejnowski</i>	5	
Using the Minimum Description Length Principle to Infer Reduced Ordered Decision Graphs	<i>Arlindo L. Oliveira and Alberto Sangiovanni-Vincentelli</i>	000
PAO Learning of One-Dimensional Patterns	<i>Paul W. Goldberg, Sally A. Goldman and Stephen D. Scott</i>	000
On-line Prediction and Conversion Strategies	<i>Nicolò Cesa-Bianchi, Yoav Freund, David P. Helmbold and Manfred K. Warmuth</i>	000

Volume 25, Nos. 2/3, November/December 1996

Special Issue on COLT '94
Guest Editor: Thomas Hancock

Guest Editor's Introduction	<i>Thomas Hancock</i>	115
The Power of Amnesia: Learning Probabilistic Automata with Variable Memory Length	<i>Dana Ron, Yoram Singer and Naftali Tishby</i>	117
CLASSIC Learning	<i>Michael Frazier and Leonard Pitt</i>	151
Rigorous Learning Curve Bounds from Statistical Mechanics	<i>David Haussler, Michael Kearns, H. Sebastian Seung and Naftali Tishby</i>	195
On the Limits of Proper Learnability of Subclasses of DNF Formulas	<i>Krishnan Pillaipakkamnatt and Vijay Raghavan</i>	237